

(19)



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 0 867 861 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
30.09.1998 Bulletin 1998/40

(51) Int. Cl.⁶: G10L 5/06

(21) Application number: 98105292.1

(22) Date of filing: 24.03.1998

(84) Designated Contracting States:
AT BE CH DE DK ES FI FR GB GR IE IT LI LU MC
NL PT SE
Designated Extension States:
AL LT LV MK RO SI

(30) Priority: 24.03.1997 US 822034

(71) Applicant:
Octel Communications Corporation
Milpitas, California 95035-7912 (US)

(72) Inventor:
The designation of the inventor has not yet been filed

(74) Representative:
Schmidt, Steffen J., Dipl.-Ing.
Wuesthoff & Wuesthoff,
Patent- und Rechtsanwälte,
Schweigerstrasse 2
81541 München (DE)

(54) Speech-responsive voice messaging system and method

(57) A method for speech-responsive voice messaging comprises the steps of generating a set of candidate results, each candidate result corresponding to a potential match between a command and an utterance, evaluating the quality of the candidate results according to a plurality of quality thresholds, and invoking one from a group of a speech user interface navigation operation and a voice messaging operation according to the quality of the candidate result evaluation.




Fig. 6

EP 0 867 861 A2

Description**BACKGROUND OF THE INVENTION****1.1 Field of the Invention**

The present invention relates to systems and methods for both voice messaging and speech recognition. More particularly, the present invention is a voice messaging system and method responsive to speech commands issued by a voice messaging subscriber.

1.2 Description of the Background Art

Voice messaging systems have become well-known in recent years. A typical Voice Messaging System (VMS) interacts with a subscriber through a Dual-Tone Multi-Frequency (DTMF), or touchtone, voice messaging User Interface (UI). During subscriber interactions, the VMS issues a voice prompt requesting the subscriber to press one or more DTMF keys to initiate corresponding operations. In the event that the subscriber presses a valid DTMF key sequence, the VMS performs a particular set of operations.

Under certain circumstances, it may be inconvenient or even dangerous for a subscriber to focus their attention on a keypad. For example, in a wireless telephone environment where a subscriber is driving or walking while on the telephone, requiring the subscriber to select an option from a set of DTMF keys could result in an accident or difficult situation. As a result, systems and methods have been developed for using speech as a means for providing hands-free interaction with a VMS, through speech-based selection of commands, user interface navigation, and entry of digits and/or digit strings.

Those skilled in the art will recognize that a conventional DTMF voice messaging UI usually has a fairly complex or extensive hierarchy of menus. Some systems that provide speech-based VMS interaction simply implement a speech UI having an identical or essentially identical menu hierarchy as a conventional DTMF UI. When a subscriber must concurrently perform multiple tasks, such as driving and VMS interaction, reducing the complexity of lower-priority tasks is very important. Thus, systems that implement a speech UI in this manner are undesirable because they fail to reduce VMS interaction complexity.

Those skilled in the art will recognize that speech recognition is an inexact technology. In contrast to DTMF signals, speech is uncontrolled and highly variable. The difficulty of recognizing speech in telephone environments is increased because telephone environments are characterized by narrow bandwidth, multiple stages of signal processing or transformation, and considerable noise levels. Wireless telephone environments in particular tend to be noisy due to high levels of background sound arising from, for example, a car

engine, nearby traffic, or voices within a crowd.

To facilitate the successful determination of a subscriber's intentions, speech-based voice messaging systems must provide a high level of error prevention and tolerance, and significantly reduce the likelihood of initiating an unintended operation. Speech-based voice messaging systems should also provide a way for subscribers to successfully complete a set of desired voice messaging tasks in the event that repeated speech recognition failures are likely. Prior art speech-based voice messaging systems are inadequate in each of these respects.

The difficulties associated with successfully recognizing subscribers' speech and determining their intentions necessitates a high level of support and maintenance to achieve optimal system performance. The availability of particular speech recognition data and system performance measures can be very useful in this regard, especially for system testing and problem analysis. Prior art systems and methods fail to provide an adequate means for flexibly controlling when and how speech recognition data and system performance measures are stored and/or generated. Moreover, prior art systems and methods fail to collect maximally useful speech recognition data, namely, the speech data generated during actual in-field system use. What is needed is a speech-responsive voice messaging system and method that overcomes the shortcomings in the prior art.

SUMMARY OF THE INVENTION

The present invention is a system and method for speech-responsive voice messaging, in which a Speech-Responsive VMS (SRVMS) preferably provides a hierarchically-simple speech UI that enables subscribers to specify mailboxes, passwords, digits, and/or digit strings. In the SRVMS, a recognition command generator and a speech and logging supervisor control the operation of a speech recognizer. A recognition results processor evaluates the quality of candidate results generated by the speech recognizer according to a set of quality thresholds that may differ on a word-by-word basis. In the preferred embodiment, the recognition results processor determines whether individual candidate results are good, questionable, or bad; and whether two or more candidate results are ambiguous due to a significant likelihood that each such result could be a valid command. The recognition results processor additionally identifies a best candidate result.

Based upon the outcome of a quality evaluation, an interpreter facilitates navigation through speech UI menus or invocation of voice messaging functions, in conjunction with a speech UI structure, a voice messaging function library, and the recognition command generator. If the recognition results processor has determined that candidate results are questionable or ambiguous, the interpreter, in conjunction with an ambi-

guity resolution UI structure and the recognition command generator, initiates confirmation operations in which the subscriber is prompted to confirm whether the best candidate result is what the subscriber intended.

In response to repeated speech recognition failures, the interpreter initiates a transfer to a DTMF UI, in conjunction with a DTMF UI structure and the voice messaging function library. Transfer to the DTMF UI is also performed in response to detection of predetermined DTMF signals issued by the subscriber while the speech UI is in context. The present invention therefore provides for both automatic and subscriber-selected transfer to a reliable backup UI.

If a best candidate result corresponds to a voice messaging function, the interpreter directs the mapping of the best candidate result to a digit sequence, and subsequently transfers control to a voice messaging function to which the digit sequence corresponds. Because the present invention provides both a speech and a DTMF UI, the mapping of candidate results allows the speech UI to seamlessly overlay portions of a standard DTMF UI, and utilize functions originally written for the DTMF UI. The present invention also relies upon this mapping to facilitate simultaneous availability of portions of the speech UI and DTMF UI while remaining within the context of the speech UI. Thus, while at particular positions or locations within the speech UI, the present invention can successfully process either speech or DTMF signals as valid input for speech UI navigation.

The SRVMS thus provides a high level of error tolerance and error prevention to successfully determine a subscriber's intentions, and further provides access to a DTMF UI in parallel with portions of the speech UI or as a backup in situations where repeated speech recognition failure is likely.

A logging unit and a reporting unit operate in parallel with the speech UI, in a manner that is transparent to subscribers. The logging unit directs the selective logging of subscriber utterances, and the reporting unit selectively generates and maintains system performance statistics on multiple detail levels.

The present Invention flexibly controls speech recognition, candidate result quality evaluation, utterance logging, and performance reporting through a plurality of parameters stored within a Speech Parameter Block (SPAB). Each SPAB preferably corresponds to a particular speech UI menu.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a block diagram of an exemplary voice messaging environment in which the present invention functions;

Figure 2 is a flowchart of a preferred minimal set of speech user-interface menu options provided to voice messaging subscribers by the present invention;

Figure 3 is a block diagram of a preferred embodiment of a Speech-Responsive Voice Messaging System constructed in accordance with the present invention;

Figure 4A is a block diagram of a preferred embodiment of a Speech Parameter Block of the present invention;

Figure 4B is a block diagram of a preferred embodiment of a vocabulary module of the present invention;

Figure 5 is a flowchart of a preferred method for providing speech-responsive voice messaging in accordance with the present invention;

Figure 6 is a flowchart of a preferred method for evaluating a speech recognition result in the present invention;

Figure 7 is a flowchart of a preferred method for confirming a speech recognition result in the present invention;

Figure 8 is a flowchart of a preferred method for utterance logging in the present invention;

Figure 9A is a graphical representation of reference times related to utterance sampling;

Figure 9B is a block diagram of a preferred utterance storage format in the present invention; and

Figure 10 is a flowchart of a preferred method for generating Customer Data Records in the present invention;

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

In the present invention, the term "subscriber" refers to a given telephone system user having direct access to voice messaging services, such as voice mail, message store and forward, and message distribution operations. The terms "nonsubscriber" and "non-user" refer to a telephone system user having no direct access to voice messaging services other than the ability to contact a subscriber, such as by entering a paging dialogue or leaving a voice message in the event that the subscriber fails to answer the nonsubscriber's call.

The terms "mobile subscriber" and "mobile nonsubscriber" are analogously defined for mobile or cellular telephone users.

Referring now to Figure 1, a block diagram of an exemplary Voice Messaging (VM) environment employing a Speech-Responsive Voice Messaging System (SRVMS) 10 is shown. In the exemplary voice messaging environment, the SRVMS 10 is coupled to a reporting system 12. Additionally, a Central Office (CO) switch 20 couples a set of subscriber telephones 30, a set of non-subscriber telephones 40, a Public-Switched Telephone Network (PSTN) 50, and the SRVMS 10. The PSTN 50 is further coupled to a Mobile Telephone Switching Office (MTSO) 70 within a cellular telephone system service area 60. The MTSO 70 exchanges infor-

mation with a set of cellular radio facilities 80 to provide telephone service to one or more mobile subscriber telephones 90 and mobile nonsubscriber telephones 92. With the exception of the SRVMS 10, the elements and their couplings shown in Figure 1 are preferably conventional.

Those skilled in the art will recognize that many variations upon the exemplary voice messaging environment of Figure 1 can be provided. For example, the MTSO 70 could be directly coupled to the CO switch 20 rather than through the PSTN 50; or the elements directed to cellular telephony could be replaced with elements representative of satellite telephony. The voice messaging environment shown in Figure 1 is useful to aid understanding, and does not limit the applicable scope of the present invention.

The SRVMS 10 provides a speech User Interface (UI) through which subscribers can verbally navigate through one or more menus to select VM service options. Those skilled in the art will understand that the provision of specific SRVMS functions may be conventionally limited to one or more particular subsets of mobile and/or non-mobile subscribers. In response to a subscriber speaking particular command words or phrases within the context of any given menu, the SRVMS 10 invokes corresponding voice messaging services.

Referring now to Figure 2, a flowchart showing a preferred minimal set of speech UI menu options provided to subscribers is shown. For each menu shown in Figure 2, the SRVMS 10 issues a voice prompt to a subscriber. Preferably, the voice prompt specifies a list of target command words or phrases, and optionally either an additional description or a voice messaging service to which each target command word or phrase corresponds. As shown in Figure 2, the preferred minimal set of speech UI menu options includes a menu for the entry of a subscriber's mailbox number; a menu for the entry of the subscriber's password; a main menu from which administrative operations or transfer to a message review menu can be selected; and the message review menu itself. Additionally, the minimal set of speech UI menu options provides submenus for skipping, canceling, or confirming particular operations. Those skilled in the art will recognize that additional menus and/or submenus, as well as menu or submenu options, can be provided. For example, a menu could be added to provide subscribers with the options of sending a message, replying to a message, or forwarding a message; or a menu could be added to support outcalling operations, in a manner readily understood by those skilled in the art. Preferably, the total number of menus and submenus through which a subscriber must navigate is kept to a reasonable number to facilitate ease of use. Exemplary voice prompts include "mailbox number please," "password please," and ""Main menu: choices are review, change greeting, change password, and hang up." Short voice prompts that convey a high level

of meaning are preferably utilized within each speech UI menu to help maximize the speed of interactions between subscribers and the SRVMS 10.

In the preferred embodiment, the speech UI is designed such that navigation through a minimum number of speech UI menus is required to access a most common set of voice messaging operations. In contrast to a standard DTMF UI, the speech UI preferably incorporates more commands into particular menus, thereby resulting in fewer menus than a DTMF UI. The preferred speech UI is therefore referred to as being hierarchically flatter than a DTMF UI. This type of speech UI enhances ease of use by reducing a subscriber's "learning curve," and aiding memorization of particular command locations within the speech UI.

The SRVMS 10 can be applied to essentially any VM environment in which verbal navigation through a speech UI may be useful. For example, the SRVMS 10 can be applied to VM environments that include essentially any wireless telephone system; or where DTMF service is unavailable, as might be the case in developing countries.

SYSTEM COMPOSITION

Referring now to Figure 3, a block diagram of a preferred embodiment of the Speech-Responsive Voice Messaging System 10 constructed in accordance with the present invention is shown. The SRVMS 10 comprises a system control unit 100, a disk and voice Input/Output (I/O) control unit 160, a data storage unit 170 upon which a database directory entry and a mailbox for each subscriber reside, at least one Digital Line Card (DLC) 180, a Telephony Interface Controller (TIC) 185 corresponding to each DLC 180, and a System Manager's Terminal (SMT) 250. The elements of the SRVMS 10 are selectively coupled via a first control bus 260 and a first data bus 262 in a conventional manner. Each TIC 185 is conventionally coupled to the CO switch 20. In the preferred embodiment, the disk and voice I/O control unit 160, the data storage unit 170, and the SMT 250 are conventional.

The system control unit 100 manages the overall operation of the SRVMS 10, in accordance with system parameter settings received via the SMT 250. The system control unit 100 preferably comprises a bus and Direct Memory Access (DMA) controller 110, a processing unit 120, and a memory 130 in which a Voice Messaging (VM) function library 132, an interpreter 134, a DTMF UI structure 136, a speech UI structure 138, and ambiguity resolution UI structure 140, a recognition command generator 142, a recognition result processor 144, a logging unit 146, a reporting unit 148, a Speech Parameter Block (SPAB) library 150, and a call statistic library 152 reside. The bus and DMA controller 110, the processing unit 120, and each element within the memory 130 is coupled via an internal bus 270. The bus and DMA controller 110 is further coupled to the first data

and control buses 260, 262, the SMT 250, as well as the reporting system 12. Preferably, the coupling maintained between the bus and DMA controller 110 and the reporting system 12 includes multiple lines, allowing data transfers according to multiple protocols.

The DLC 180 exchanges voice data with the CO switch 20, processes DTMF signals, and performs speech recognition and logging operations under the direction of the system control unit 100. The DLC 180 preferably comprises a DLC bus controller 190, a DLC processing unit 200, a Coder/Decoder (CODEC) 210, and a DLC memory 220. A speech recognizer 222, a DTMF processor 224, a template library 226, a logging buffer 228, a speech and logging supervisor 230, a phrase expander 232, an auto-response library 234, a Pulse Code Modulation (PCM) data buffer 236, and a signal conditioner 238 reside within the DLC memory 220. Each element within the DLC memory 220 is coupled to the DLC bus controller 190 and the DLC processing unit 200 via a second data bus 280. The DLC bus controller 190 is coupled to the DLC processing unit 200 via a second control bus 282. Additionally, the DLC bus controller 190 is coupled to the first data and control buses 260, 262. The CODEC 210, the phrase expander 232, the signal conditioner 238, and the DTMF processor 224 are preferably conventional. The CODEC 210 is coupled to the PCM data buffer 236 and the DLC bus controller 190 to effect DMA-type operations between the PCM data buffer 236 and the TIC 185.

KEY REQUIREMENTS AND FUNCTIONAL ABILITIES

In order to provide successful speech-responsive VM, several key interrelated requirements must be met. The nature of these key requirements and the manner in which they are facilitated by individual elements within the SRVMS 10 is hereafter described.

I. A first key requirement is the ability to detect a subscriber's utterance, and identify particular command words or phrases to which the utterance may correspond. This ability is provided by the speech recognizer 222 in conjunction with the template library 226 and autoresponse library 234.

The speech recognizer 222 is preferably conventional, and provides speaker-independent recognition of subscriber utterances in a discrete recognition mode when detection of command words and/or individual digits is required, or a continuous recognition mode when detection of digit strings is required. The speech recognizer 222 also preferably provides a connected recognition mode in which detection of particular conditions results in an automatic restart of a recognition attempt, as described in detail below. When in continuous recognition mode, the speech recognizer 222 can preprocess an utterance to facilitate the identification

of individual digits. In the preferred embodiment, the speech recognizer 222 can additionally provide speaker-dependent or speaker adaptive speech recognition.

The template library 226 stores word templates and corresponding word identifications (IDs), which define each valid command word within the speech UI for the speech recognizer 222 in a manner those skilled in the art will readily understand. The autoresponse library 234 stores word templates and corresponding word IDs that define autoresponse command words that the speech and logging supervisor 230 can independently act upon, as described in detail below.

The speech recognizer 222 initiates a recognition attempt under the direction of the speech and logging supervisor 230, as described in detail below. During a recognition attempt, the speech recognizer 222 attempts to determine the closest match or matches between a subscriber's utterance and a vocabulary. Herein, a vocabulary is defined as a subset of the word templates stored in the template library 226. The vocabulary corresponds to the command words or phrases available within a particular speech UI menu. Thus, a vocabulary is an organization of particular word templates. Upon completion of a recognition attempt, the speech recognizer 222 returns recognition results to the speech and logging supervisor 230. Preferably, the recognition results comprise a set of candidate results, where each candidate result includes a candidate word ID and at least one score corresponding to each candidate word ID. To aid understanding, the description herein assumes a single score is associated with each candidate word ID. Predetermined candidate results are preferably reserved for indicating the occurrence of a timeout condition, an Out-of Vocabulary Word (OVW), an unresolvable error, or other "match not possible" conditions.

A variety of recognizer parameters control the manner in which the speech recognizer 222 operates. In the preferred embodiment, the following can be specified by the recognizer parameters: type of recognition to be performed; timeout information; a minimum and a maximum acceptable string length; a reference to a particular vocabulary; a number of candidate results required; and score control information.

II. A second key requirement is the ability to issue appropriately-structured commands for controlling the speech recognizer 222. This is facilitated through the recognition command generator 142 and the speech and logging supervisor 230. In response to a call issued by the interpreter 134, the recognition command generator 142 issues a recognition parameter directive to the speech and logging supervisor 230. Preferably, the recognition

parameter directive specifies the previously described recognizer parameters. In response to the recognition parameter directive, the speech and logging supervisor 230 initializes the speech recognizer 222.

The recognition command generator 142 additionally issues a recognition request to the speech and logging supervisor 230. Upon receiving the recognition request, the speech and logging supervisor 230 directs the speech recognizer 222 to initiate a recognition attempt. The speech and logging supervisor 230 additionally initiates the operation of the DTMF processor 224 such that the occurrence of a particular DTMF signal or a hang-up condition can be detected.

After the speech recognizer 222 generates a set of candidate results (or after the DTMF processor 224 generates a result), the speech and logging supervisor 230 either performs autoresponse operations, or transfers the candidate result sets (or a DTMF Signal ID) to the control unit memory 130 and returns a value to the interpreter 134 to initiate result processing operations. The sequence of events beginning with the recognition command generator's issuance of the recognition request and ending with the return of a value to the interpreter 134 is referred to herein as a recognition event.

In the preferred embodiment, the speech and logging supervisor 230 performs autoresponse operations. In the event that the speech recognizer 222 has detected a particular autoresponse command word stored in the autoresponse library 234. Preferably, the autoresponse words include "faster," "slower," "louder," and "softer." The speech and logging supervisor 230 performs a set of operations corresponding to the detected autoresponse command word. Detection of "faster" or "slower" results in faster or slower message playback, respectively; and detection of "louder" or "softer" respectively results in a volume increase or decrease. The speech and logging supervisor 230 can also perform autoresponse operations in response to the detection of particular error conditions. Autoresponse operations are preferably enabled via a connected recognition mode. After performing autoresponse operations, the speech and logging supervisor 230 initiates another recognition attempt in accordance with the most-recent recognition request.

III. A third key requirement for providing successful speech-responsive VM is the ability to analyze or evaluate the quality of the candidate results. This ability is facilitated through the recognition result processor 144. Following the completion of a recognition event, the recognition result processor 144 determines the whether candidate results are good, bad or questionable. The detailed operations performed by the recognition result processor 144 are

described below with reference to Figure 6.

IV. A fourth key requirement for providing successful speech-responsive VM is the ability to control which portion of the speech UI is presented to the subscriber at any point in time, and selectively transition from one portion of the speech UI to another or invoke a voice messaging function based upon the outcome of the evaluation performed by the recognition result processor 144. This ability is facilitated through the interpreter 134, the speech UI structure 138, and the VM function library 132.

In the preferred embodiment, each UI structure 136, 138, 140 comprises a data structure that hierarchically organizes references to sequences of program instructions that implement either UI navigation operations or VM functions. Each such program instruction sequence is preferably stored within the VM function library 132. The aforementioned hierarchical organization corresponds to the menus and submenus available to subscribers. In the preferred embodiment, each UI structure 136, 138, 140 comprises a tree.

For implementing the speech UI, the interpreter 134 selects or maintains a reference to a position or location within the speech UI structure 136. Based upon the current location within the speech UI structure 136, a value returned by the speech and logging supervisor 230, and the outcome of the recognition result processor's candidate result set evaluation, the interpreter 134 directs control transfers to appropriate program instruction sequences within the VM function library 132. In the preferred embodiment, the interpreter 134 initiates control transfers via event-driven case-type statements. A recognition event that culminates in the execution of a VM function is referred to herein as a communication.

In the present invention, a particular UI is implemented using the interpreter 134, a given UI structure 136, 138, 140, and the set of program instruction sequences within the VM function library 132 that are referenced by the given UI structure 136, 138, 140. Thus, the speech UI structure 136, the interpreter 134, and a particular group of VM functions together implement the present invention's speech UI. Similarly, the DTMF UI structure 136 in conjunction with the interpreter 134 and VM function library 132, implement a DTMF UI, which in the preferred embodiment is defined in accordance with Voice Messaging User Interface Forum (VMUIF) standards. The ambiguity resolution UI structure 140, along with the interpreter 134 and portions of the VM function library 132, implement a confirmation menu within the speech UI, through which a subscriber is prompted to confirm a previous response, as described in detail below with reference to Figure 7.

Those skilled in the art will recognize that each

UI is implemented in accordance with threaded code techniques, in particular, threaded code techniques as commonly defined in the context a programming language such as Forth or Java. While any given UI could be implemented in another manner as readily understood by those skilled in the art, the implementation of a UI as described above results in enhanced portability across different system types, fast operation, reduced storage requirements, and also facilitates simpler system development and testing.

V. A fifth key requirement for providing successful speech-responsive VM is the ability to have one or more portions of a DTMF UI available in parallel with the speech UI, as well as the ability to rely upon the DTMF UI as a backup in situations where repeated speech recognition failures seem likely. This requirement is satisfied by the interpreter 134, the VM function library 132, and the DTMF UI structure 136. Having portions of the DTMF UI available in parallel with the speech UI facilitates the processing of subscriber input regardless of whether such input is in the form of speech or DTMF signals. This concurrent UI provision provides for a) situations in which it may be desirable to process either speech or DTMF signals, and remain within the context of the speech UI, such as when subscriber entry of a mailbox number or password is required; and b) the transfer out of the speech UI and into the DTMF UI in response to receipt of particular DTMF input.

The presence of the DTMF UI to serve as a backup to the speech UI makes the SRVMS 10 more reliable than systems in which speech is the sole input means for UI navigation. In situations where speech recognition is consistently problematic, the DTMF UI enables subscribers to successfully complete their VM tasks.

Those skilled in the art will recognize that transfer to the DTMF UI is only viable in telephony environments in which DTMF is available, unless rotary dialing detection and mapping functionality available for mapping rotary signals to DTMF. Such functionality could be provided, for example, by hardware and/or software residing upon the line card 180. Those skilled in the art will recognize that providing a speech UI in a non-DTMF environment may be desirable because the entry of information by rotary dialing can be quite time consuming.

VI. A sixth key requirement for providing successful speech-responsive VM is the ability to control the issuance of selectively-interruptable prompts and messages to the subscriber. This is facilitated through the interpreter 134, a UI structure 136, 138, 140, at least one VM function within the VM function library 132, and the phrase expander 232. In the preferred embodiment, the phrase expander 232 is responsive to signals issued by the DTMF proces-

sor 224 and the speech and logging supervisor 230, and will play a prompt or message until a DTMF signal has been detected or the speech and logging supervisor 230 returns recognition results to the recognition result processor 144. Additionally, a prompt may be halted at an earlier time, when the speech recognizer 222 detects the beginning of a recognizable utterance (such as the start of a digit string). This capability is referred to herein as "barge-in," and is selectively performed in accordance with a set of interruption codes. Providing for voice prompt and message interruptability helps maximize the speed of interactions between the subscriber and the SRVMS 10. In the preferred embodiment, recognition results are not returned to the recognition result processor 144 after autoresponse operations, and hence a prompt will continue playing during and after autoresponse operations.

VII. A seventh key requirement for providing successful speech-responsive VM is the ability to selectively generate and analyze SRVMS performance information. This is facilitated by the logging unit 146, the reporting unit 148, and the speech and logging supervisor 230. The generation and analysis of SRVMS performance information is particularly useful for identifying problems, and tracking the manners in which the system is used. The detailed operations performed by the logging unit 146 and the reporting unit 148 are described below with reference to Figures 8 through 10.

CONTROL PARAMETERS

The present invention relies upon a variety of parameters for controlling the initiation, evaluation, logging, and reporting of speech recognition events. For each menu within the speech UI, a corresponding SPAB 300 within the SPAB library 150 stores these parameters. Referring now to Figure 4A, a block diagram of a preferred embodiment of a SPAB 300 is shown. Each SPAB 300 is preferably a data structure that comprises a first data field 302 for storing a list of logging and reporting parameters; a second data field 304 for storing a list of speech recognition control parameters, as well as the previously mentioned interruption codes; a third data field 306 for storing a list of quality thresholds, which are described in detail below; a fourth data field 308 for storing a digit mapping list 308, which is used for mapping word IDs to voice messaging functions, as described in detail below; and a fifth data field 310 for storing a list of references to vocabulary modules.

The logging parameters specify the manners in which the logging unit 146 directs the logging of subscriber utterances, and preferably include condition codes that selectively specify the following:

- whether logging shall be pseudo-random at a call-

level, communication level, or recognition event level, selectable in terms of a particular number per 1000 calls, communications, or recognition events, respectively (ranging from 0 per 1000 for never, to 1000 per 1000 for always);

- SRVMS port number;
- one or more subscriber mailboxes;
- one or more menus within the speech UI;
- specific word IDs;
- recognition types for which logging is to occur;
- whether to log good recognitions;
- whether to log bad recognitions;
- whether to log questionable words or confusing word pairs;
- whether to log commands or digits;
- specific error or OVW conditions to be logged; and
- sampling parameters.

The sampling parameters are used by the speech and logging supervisor 230, and preferably specify whether logging is to occur for raw or preprocessed (i.e., echo-canceled) speech; and timing definitions that indicate at what point during speech recognition logging is to begin and end. The detailed operations performed by the logging unit 146 are described below with reference to Figures 8, 9A, and 9B.

The reporting parameters control the manner in which the reporting unit 148 operates, and preferably specify whether reporting is to occur, plus control conditions indicating whether reporting is to be performed on a per-recognition basis, a per-communication basis, or a per-call basis. The detailed operations performed by the reporting unit 148 are described below with reference to Figure 10.

The speech recognition parameters specify initialization and recognition settings for the speech recognizer. In the preferred embodiment, the speech recognition parameters indicate a type of recognition to be performed; timeout information; a minimum and a maximum acceptable string length; a reference to a particular vocabulary; a number of candidate results required; score control information; and error control information.

Referring also now to Figure 4B, a block diagram of a preferred embodiment of a vocabulary module 320 is shown. Each vocabulary module 320 is a data structure comprising a first data field 322 for storing a list of word IDs, and a second data field 324 for storing a word or phrase corresponding to each word ID. Any given vocabulary module 320 specifies the command words or phrases that are available to the subscriber within a particular menu of the speech UI. In the preferred embodiment, a collection of SPABs 300 exist for each language supported by the SRVMS 10.

EMBODIMENT DETAILS

In the preferred embodiment, each of the inter-

preter 134, the recognition command generator 142, the recognition result processor 144, the logging unit 146, and the reporting unit 148 comprise a sequence of program instruction sequences that are executable by the processing unit 120 and stored in the memory 130. Similarly, each of the speech recognizer 222, the DTMF processor 224, the speech and logging supervisor 230, the phrase expander 232, and the signal conditioner 238 comprise program instruction sequences executable by the DLC processing unit 200 and stored in the DLC memory 220. The DLC processing unit 200 is preferably implemented using a commercially-available Digital Signal Processor (DSP). Those skilled in the art will recognize that one or more portions of the aforementioned elements may instead be implemented as hardware in an alternate embodiment, and will also understand that the DLC processing unit 200 does not have to be a DSP (for example, a Pentium processor (Intel Corporation, Santa Clara, CA) could be used).

In an exemplary embodiment, the SRVMS 10 is an Octel Sierra system (Octel Communications Corporation, Milpitas, CA) having the elements shown within the system controller memory 130 and the DLC memory 220; an 80486 microprocessor (Intel Corporation, Santa Clara, CA) serving as the DLC bus controller 190; a Texas Instruments C31 DSP (Texas Instruments Corporation, Dallas, TX); Portable Recognizer Library (PRL) software (Voice Processing Corporation, Cambridge, MA); and a personal computer having a Pentium or similar processor to serve as the SMT 250, which is coupled to the bus and DMA controller 110 via a conventional X.25 coupling and a Small Computer System Interface (SCSI) bus. In an alternate embodiment, the SRVMS 10 could be implemented in a unified or integrated voice messaging system, such as that described in U.S. Patent No. 5,557,659, entitled "Electronic Mail System Having Integrated Voice Messages." In such implementations, elements of the SRVMS 10 shown in Figure 3 reside within a voice server coupled to an electronic mail system, in a manner readily understood by those skilled in the art.

Those skilled in the art will additionally recognize that in yet another embodiment, the SRVMS 10 could be implemented in a single-processor system. In such an embodiment, the DLC processing unit 200 is not present (or equivalently, the DLC processing unit 200 and the processing unit 120 are one and the same), and elements 222, 224, 226, 228, 230, 232, 234, 236, 238 within the DLC memory 220 of Figure 1 are instead implemented within the control unit memory 130, with the exception of the CODEC 210 in the event that DMA-type transfers from the TIC 185 are required.

DETAILED OPERATION

The manner in which the aforementioned system elements interact sequentially and/or in parallel to implement speech-responsive VM in an essentially

seamless manner is described in detail hereafter with reference to Figures 5 through 10.

Referring now to Figure 5, a flowchart of a preferred method for providing speech-responsive voice messaging in accordance with the present invention is shown. In the preferred embodiment, the operations performed in Figure 5 are initiated when the interpreter 134 transfers a reference to an initial SPAB 300 to the recognition command generator 142 in response to an incoming call notification received from the DLC bus controller 190.

The preferred method begins in step 500 with the recognition command generator 142 selecting the initial SPAB 300 for consideration. In the preferred embodiment, the first SPAB 300 corresponds to a mailbox number entry menu. Those skilled in the art will recognize that the first SPAB 300 could correspond to some other menu, such as a welcome menu that could facilitate offering a subscriber a choice between use of the speech and DTMF UIs. Next, the recognition command generator 142 retrieves the recognizer parameters within the currently-selected SPAB 300 in step 502, and issues a recognition parameter directive to the speech and logging supervisor 230 step 504. The speech and logging supervisor 230 subsequently initializes the speech recognizer 222 accordingly. Then, in step 506, the recognition command generator 142 issues a recognition request, thereby initiating a recognition event. After step 506, a voice messaging function within the VM function library 132 selects a current prompt, and issues a prompt notification to the phrase expander 232 in step 508. In turn, the phrase expander 232 issues the current prompt to the subscriber in a conventional manner, that is, via the PCM data buffer 236, the CODEC 210, and the TIC 185. The prompt is preferably played until a DTMF signal has been detected, or the speech and logging supervisor 230 returns a candidate result set to the control unit 100.

Following step 508, the recognition result processor 144 retrieves the candidate result set in step 510. In the preferred embodiment, the interpreter 134 initiates control transfer to the recognition result processor 144 in response to the speech and logging supervisor's return of a value indicating a candidate result set requires evaluation. The recognition result processor 144 subsequently evaluates the quality of the returned candidate results in step 512, as described in detail below with reference to Figure 6, and preferably returns a value to the interpreter 134 that indicates the outcome of this evaluation.

Based upon the value received from the recognition result processor 144, the interpreter 134 determines whether recognition is to be repeated in step 514. If the outcome of the recognition result processor's evaluation indicates that the subscriber's response was bad, and a recognition repeat count has not been exceeded, recognition must be repeated. A bad response could result from any significant audible event that was not an

expected word, possibly arising from, for example, excessive background sound. In the event that recognition must be repeated, the preferred method returns to step 506 to initiate another recognition event. In the preferred embodiment, the current prompt issued in step 508 can vary according to the number of times recognition has been repeated.

In the event that a subscriber's response was bad and the repeat count has been exceeded, the interpreter 134 transitions to the DTMF UI via steps 516 and 518. After step 518, the preferred method ends.

If neither recognition repetition nor transfer to the DTMF UI are required, the interpreter 134 determines whether recognition confirmation is required in step 520. In the present invention, confirmation is required when the outcome of the evaluation indicates a questionable or ambiguous response. If confirmation is required, the interpreter 134 selects a position or location within the ambiguity resolution UI structure 140, and transfers a reference to a confirmation SPAB 300 to the recognition command generator 142 in step 522 to initiate confirmation operations as described in detail below with reference to Figure 7. After step 522, the interpreter determines whether the confirmation was successful in step 524. If not, the preferred method returns to step 506.

When confirmation is not required in step 520, or after a successful confirmation in step 524, the interpreter 134 transfers control to a mapping function that maps the best candidate word ID to a digit sequence in step 530. The mapping function relies upon data within the current SPAB 300 to perform mapping operations. The interpreter 134 subsequently determines whether the mapped digit sequence corresponds to a speech UI navigation operation in step 532. If so, the interpreter 134 selects a position or location within the speech UI in step 534. In the event that a VM action is required rather than speech UI navigation, the interpreter transfers control to a VM function that corresponds to the mapped digit sequence in step 536. In the preferred embodiment, a digit string is interpreted as a single entity.

A VM function that directs message playback preferably operates in conjunction with the recognition command generator 142 and recognition result processor 144 such that the recognition and evaluation of subscriber utterances is selectively performed while a message is played to the subscriber. This in turn helps maximize interaction speed between the SRVMS 10 and the subscriber.

After step 536, the interpreter 134 updates a communication count in step 538. Herein, a communication is defined as a successful speech UI interaction with a subscriber that culminates in the execution of a voice messaging function. The communication count is selectively utilized by the reporting unit 148, as described in detail below with reference to Figure 10.

Those skilled in the art will recognize that the digit sequence generated in step 530 could correspond to a

sequence of DTMF commands that would request the same voice messaging service had the subscriber interacted with the DTMF UI. Thus, the mapping performed in step 530 allows the SRVMS 10 to directly use one or more program instruction sequences originally written for implementing voice messaging functions on a DTMF-only system. Thus, the speech UI provided by the present invention can partially or entirely overlay a conventional DTMF UI, potentially increasing system reliability and/or reducing system development time. Moreover, the speech UI provided by the present invention can seamlessly overlay two or more non-identical DTMF UIs.

After steps 534 or 538, the interpreter 134 determines whether the subscriber's call is complete in step 540. If so, the preferred method ends. Otherwise, the preferred method proceeds to step 500, where the interpreter 134 selects an appropriate SPAB 500 for consideration. Call completion is preferably indicated when the DTMF processor 224 detects a hangup condition, or a command word or phrase such as "hang up" is successfully detected and processed.

In the preferred embodiment, successful recognition of the word "help" causes the interpreter 134 to transition to a particular help menu within the speech UI via the selection of a corresponding help SPAB 300. Preferably, a variety of help SPABs 300 exist, to facilitate the implementation of context-sensitive user assistance from any main speech UI menu. The interpreter's selection of a particular help SPAB 300 is thus based upon the position or location within the speech UI from which the subscriber requested help.

Referring now to Figure 6, a flowchart of a preferred method for evaluating a speech recognition result (step 510 of Figure 5, and step 710 of Figure 7) is shown. The preferred method begins in step 600 with the recognition result processor 144 determining whether the candidate result set indicates that an unrecoverable error or a timeout condition had occurred. If so, the recognition result processor 144 sets a bad result status indicator in step 604, and increments a repeat count in step 606. When evaluating the quality of confirmation results, the recognition result processor 144 increments a confirmation repeat count; otherwise, the recognition result processor 144 increments a recognition repeat count. If an appropriate repeat count limit has been exceeded, the recognition result processor 144 sets a corresponding limit exceeded status via steps 606 and 610. In the event that the appropriate repeat count limit has not been exceeded, the recognition result processor 144 sets a repeat status indicator in step 608. After either of steps 608 or 610, the preferred method ends.

If no error or timeout occurred, the recognition result processor 144 selects a first candidate result in step 620. The recognition result processor 144 then compares the score within the selected candidate result with a group of threshold scores corresponding to the selected candidate result's word ID in step 622. Prefer-

ably, the threshold scores for each valid word ID within a speech UI menu stored are stored in the current SPAB 300. In the preferred embodiment, a first threshold score establishes a first quality level above which the candidate result is deemed "good." A second threshold score establishes a second quality level, below which the candidate result is deemed "bad." Between the first and second quality levels, the candidate result is deemed "questionable." Those skilled in the art will recognize that in an alternate embodiment, additional quality threshold levels could be defined, such as "very good." Those skilled in the art will also recognize that in embodiments where the speech recognizer returns multiple types of scores for a single candidate word ID, separate types of threshold scores could be analogously defined. In an alternate embodiment, the recognition result processor 144 additionally performs statistical language modeling operations to aid quality evaluation.

After step 622, the recognition result processor 144 marks the currently-selected candidate result in accordance with its quality designation in step 624. The recognition result processor 144 then determines whether another candidate result requires consideration in step 626. If so, the preferred method returns to step 620.

Once each candidate result has been considered, the recognition result processor 144 determines whether at least one candidate result has been designated as "good" in step 630. If so, the recognition result processor 144 determines whether multiple good candidate results are present in step 632. If only one candidate result has been designated as good, the recognition result processor 144 sets a good result status indicator in step 638, and returns this candidate result in step 644, after which the preferred method ends.

When multiple good candidate results are present, the recognition result processor 144 examines the score differences between each good candidate result in step 634, and determines whether a minimum score difference threshold is exceeded in step 636. If the minimum score difference threshold is exceeded, the recognition result processor 144 sets the good result status indicator in step 628, and returns the best candidate result in step 644, after which the preferred method ends. In the preferred embodiment, the best candidate result is defined as the least-uncertain good candidate result (as indicated by the score associated with the word ID), provided the minimum score difference threshold is exceeded. If the minimum score difference threshold is not exceeded, the recognition result processor 144 returns a confirmation required status indicator in step 642, after which the preferred method proceeds to step 644. Thus, the present invention ensures that the generation of potentially ambiguous yet good recognition results in asking the subscriber for confirmation.

In the event that a good candidate result is not present in step 630, the recognition result processor

144 determines whether a questionable candidate result is present in step 640. If so, the preferred method proceeds to step 642. Otherwise, the preferred method proceeds to step 602. In the preferred embodiment, the recognition results processor 144 evaluates candidate results expected to correspond to digit strings such that the quality or validity of any given number within the string is determined.

Referring now to Figure 7, a flowchart of a preferred method for confirming a speech recognition result (step 540 of Figure 5) is shown. The preferred method begins in step 700 with the recognition command generator 142 retrieving data within the confirmation SPAB 300. Next, the recognition command generator 142 issues a recognition parameter directive to the speech and logging supervisor 230, which sets recognizer parameters as indicated in the confirmation SPAB 300.

A voice messaging function then determines a current confirmation prompt, and issues a confirmation prompt notification to the phrase expander 232 in step 704. Preferably, the confirmation prompt notification includes a reference to the current confirmation prompt, plus the word ID of the word or phrase requiring confirmation, such that the subscriber is presented with the best word or phrase candidate during the prompt. For example, if the word "review" required confirmation, the current confirmation prompt plus the word ID in question would be presented to the subscriber in a manner such as "Did you say review? Please answer yes or no." The phrase expander 232 issues the current confirmation prompt and the word under consideration to the subscriber in a manner readily understood by those skilled in the art. In the preferred embodiment, interruption of a confirmation prompt is not allowed.

Following step 704, the recognition command generator 142 issues a recognition request, thereby initiating a recognition event in step 706. The speech and logging supervisor 230 preferably returns candidate results for the confirmation to the control unit memory 130, and returns a value to the interpreter indicating quality evaluation is required. The interpreter 134 transfers control to the recognition result processor 144 in response.

In steps 708 and 710, the recognition result processor 144 respectively retrieves and evaluates the candidate results returned after the subscriber was prompted for confirmation. Step 710 is performed in accordance with the description of Figure 6 above. Next, in step 712, the interpreter 134 determines whether the confirmation result was good in step 712. If so, interpreter 134 sets a successful confirmation status indicator in step 718, after which the recognition command generator 142 restores the recognizer parameters specified within the previously-selected SPAB 300 (i.e., the SPAB 300 most-recently selected via step 500 of Figure 5) in step 720. After step 720, the preferred method ends.

If the result of the confirmation was not good, the interpreter determines whether reconfirmation is

required in step 714. Reconfirmation is preferably called for when the recognition result processor 144 has set either the confirmation required status or the repeat status indicator. When reconfirmation is required, the method preferably returns to step 706 to initiate another recognition event. In the preferred embodiment, the current confirmation prompt issued in step 704 varies according to the number of reconfirmation attempts made.

10 If the recognition result processor 144 has set the confirmation limit exceeded status, the interpreter 134 determines that reconfirmation is not required in step 714, and sets an unsuccessful confirmation status indicator in step 716. After step 716, the preferred method proceeds to step 720.

UTTERANCE LOGGING

Recording or logging of subscriber utterances is 20 highly useful for aiding system testing and verification, periodic vocabulary building, and problem analysis. Utterance logging, however, requires significant amounts of storage, and thus logging can be quite costly. In the present invention, the logging unit 146 and 25 the speech and logging supervisor 230 control the selective logging of subscriber utterances in accordance with the logging parameters specified in each SPAB 300, such that logging costs can be minimized.

Referring now to Figure 8, a flowchart of a preferred 30 method for utterance logging in the present invention is shown. In the preferred embodiment, the logging unit 146 operates transparently during a call, monitoring the operation of the interpreter 134, the recognition command generator 142, and the recognition result processor 144. The preferred method begins in step 800 with the logging unit 146 examining the logging parameters within the currently-selected SPAB 300 (i.e., the SPAB 300 selected in step 500 of Figure 5) to determine whether utterance logging is required during the current 35 call. If not, the preferred method ends.

If utterance logging is required, the logging controller 146 establishes the current logging conditions in 40 accordance with the logging parameters in step 802. In the preferred embodiment, the logging parameters indicate various conditions under which logging is required, as previously specified in relation to Figure 4A. The logging unit 146 next determines in step 804 whether the 45 next recognition event is to be logged. If so, the logging unit 146 issues a set of sampling parameters to the speech and logging supervisor 230 in step 806. The sampling parameters preferably specify whether utterance logging is to begin according to the following reference time definitions:

50

- 55 • at the start of a recognition attempt;
- when an audio signal has been detected that has a volume and spectral composition that suggests speech, defined herein as the "start of speech"; and

- when the speech recognizer 222 is confident that an utterance is meaningful, and has started template matching processes, defined herein as the "start of utterance."

The sampling parameters additionally specify whether utterance logging is to end according to the following reference time definitions:

- after a predetermined time has elapsed since the start of utterance logging;
- after an end to speech-like data has been detected, defined herein as "end of speech"; and
- following the generation of candidate results, defined herein as "end of utterance."

Referring also now to Figure 9A, a graphical representation of the reference times defined above is shown. To compensate for time delays in assessing the aforementioned reference times, a buffer preferably holds audio data corresponding to the most-recent 1 second interval, such that the sampling period can be extended approximately 0.5 to 1 second relative to the start and end times shown. The speech and logging supervisor 230 directs logging during recognition attempts, and stores logged utterances in the logging buffer 228.

Referring again to Figure 8, after step 806, the logging unit 146 determines whether the recognition result processor 144 has completed the quality evaluation for the current candidate results in step 808. If not, the preferred method remains at step 808. Once the final result of the most recent recognition event is known, the logging unit 146 determines whether any criteria specified in the logging parameters are matched in step 810. If so, the logging unit 146 instructs the speech and logging supervisor 230 to save an utterance header and the utterance recorded during the most recent recognition event in step 812. The utterance header preferably includes a reference to a position or location within the speech UI; a retry count; a communication count; the candidate result set generated by the speech recognizer 222; timing data issued by the recognizer; timing data related to prompt playing and interruption; and timing data corresponding to the arrival of external events such as a DTMF signal or a hang-up. The utterance itself is preferably encoded according to 8-bit mu-law protocols. Each utterance header and corresponding utterance is preferably saved in the logging buffer 228, at a location given by a reference or pointer to an available storage location within the logging buffer 228. The logging unit 146 preferably maintains this pointer. Upon completion of step 812, the logging unit 146 examines the current logging parameters and determines whether the saved utterance should be retained for later use in step 816. Under particular circumstances, knowledge of whether logged utterances should be saved cannot be ascertained until the subscriber's call has proceeded to

a certain point within the speech UI. For example, the initiation of logging preferably occurs at the beginning of a call. If logging is to occur for a particular password number, however, the subscriber's password number will not be known until the call has progressed to the point at which the subscriber's utterance(s) made within the context of the password entry menu have been successfully recognized and processed.

If the utterance data is to be retained, the logging unit 146 updates the logging buffer storage location reference to a next available location in step 816.

After step 816, or after steps 804, 810, and 814, the logging unit 146 determines whether the current call is complete in step 818. If not, the preferred method returns to step 804. After the current call is complete, the logging unit 146 generates call header information in step 820, and subsequently transfers the call header information and the set of saved utterances to either the data storage unit 170 or the reporting system in step 822. In the preferred embodiment, the call header information comprises a mailbox ID, a time stamp, and possibly a reference to a Customer Data Record (CDR), which is described in detail below with reference to Figure 10. After step 822, the preferred method ends.

Referring also now to Figure 9B, a block diagram of a preferred utterance storage format 900 is shown. In the preferred utterance storage format, a call header 902 is followed by utterance header/utterance audio data sequences 904. Within the logging buffer 228, a pointer to a current logging location and a previous logging location are maintained in a manner readily understood by those skilled in the art.

REPORTING

The generation of system performance data is highly useful for system problem analysis. In the present invention, the reporting unit 148 selectively generates various Customer Data Records (CDRs), which store particular system performance statistics. In the preferred embodiment, the reporting unit 148 operates transparently during a call, monitoring the operation of the interpreter 134, the recognition command generator 142, and the recognition result processor 144 to track system performance and generate CDRs in accordance with the reporting parameters specified in each SPAB 300. For the generation of each CDR, the reporting unit 148 maintains a set of statistics within the call statistic library 152.

In the preferred embodiment, the reporting unit 148 selectively generates a recognition-level CDR, a communication-level CDR, a call-level CDR, and/or a scary-level CDR. The recognition-level CDR preferably specifies the following:

- the results of each recognition within a communication;
- the response of the system to predetermined rec-

ognition results, as specified within the current SPAB 300;

- a logging status for each recognition;
- duration of each recognition event; and
- candidate word IDs and corresponding scores for each recognition event.

The following are preferably specified by the communication-level CDR:

- a result indicating an action taken following a communication;
- the proportion of subscriber inputs requiring predetermined numbers of recognition attempts, where the predetermined numbers are specified by the SMT 250;
- the number of incorrect attempts;
- the number of timeouts;
- whether an affirmative confirmation occurred; and
- time duration of the communication.

The call-level CDR provides the following information:

- the proportion of a call in which speech was used;
- the proportion of digit strings in which speech was used;
- the proportion of digit string inputs requiring predetermined numbers of recognition attempts;
- the proportion of recognition events in which a timeout occurred;
- the proportion of recognition events requiring confirmation;
- the proportion of recognition events that failed;
- average duration of recognition events; and
- average communication duration.

Finally, the summary-level CDR contains the following information:

- the proportion of calls in which subscribers reverted to using DTMF;
- the proportion of calls in which the SRVMS 10 reverted to the DTMF UI;
- the proportion of calls in which the speech UI was re-invoked; and
- the proportion of calls in which a hang-up condition followed an unsuccessful recognition;

Those skilled in the art will readily understand the manner in which the aforementioned information can be generated and/or updated by tracking the operations performed by the interpreter 134, the recognition command generator 142, the recognition result processor 144, the speech and logging supervisor 230, and the speech recognizer 222. Those skilled in the art will additionally recognize that additional or fewer statistics could be generated in an alternate embodiment, according to the usefulness of particular information.

Referring now to Figure 10, a flowchart of a preferred method for creating Customer Data Records is shown. The preferred method begins in step 1000 with the reporting unit 148 retrieving the reporting parameters specified within the current SPAB 300 to establish current reporting conditions. Next, the reporting unit 148 determines whether a recognition-level CDR is to be generated in step 1002. If so, the reporting unit 148 monitors recognition results and recognition result evaluation processes, and generates and/or updates recognition statistics in steps 1004 and 1006.

After step 1002 or step 1006, the reporting unit 148 determines whether the current communication is complete in step 1008. If not, the preferred method returns to step 1002. Once the current communication is complete, the reporting unit 148 determines whether generation of a communication-level CDR is required in step 1010. If so, the reporting unit 148 generates and/or updates communication statistics in step 1012. After step 1010 or step 1012, the reporting unit 148 determines whether the current call is complete in step 1014. If not, the preferred method returns to step 1002.

Upon completion of the current call, the reporting unit 148 determines whether a call-level CDR should be generated, and, if so, generates and/or updates call statistics in steps 1016 and 1018, respectively. If call-level CDR generation is not required, or after step 1018, the reporting unit 148 generates each required CDR in step 1020, using the statistics maintained in the call statistic library 152. Preferably, each CDR comprises a data file in which the appropriate statistical information resides. After the CDRs have been generated, the reporting unit 148 directs their transfer to the reporting system 12.

While the present invention has been described with reference to certain preferred embodiments, those skilled in the art will recognize that various modifications can be provided. For example, speaker-dependent recognition could be employed to substitute a subscriber-generated keyword with a corresponding string of digits.

This and other variations upon the present invention are provided within the context of the embodiments described herein, which are limited only by the following claims.

45 Claims

1. A method for speech-responsive voice messaging comprising the steps of:

- 50 a) generating a set of candidate results, each candidate result corresponding to a potential match between a command and an utterance;
- b) evaluating the quality of the candidate results according to a plurality of quality thresholds; and
- 55 c) invoking one from a group of a speech user interface navigation operation and a voice messaging operation according to the quality of the

candidate result evaluation.

2. The method according to claim 1, which includes also Dual-Tone Multi-Frequency (DTMF) signal-responsive voice messaging. 5

3. The method according to claim 1 or 2, wherein:

step a) comprises the detection of an utterance and the identification of particular command words or phrases to which the utterance may correspond, wherein the steps of detection and/or identification are carried out by a speech recognizer (222), preferably in conjunction with a template library (226) and an autore-sponse library (234), and wherein the speech recognizer (222) is preferably controlled by a recognition command generator (142) and a speech and logging supervisor (230). 10 15

4. The method according to any of claims 1 to 3, wherein 20

step c) comprises controlling which portion of a speech user interface is presented to a subscriber, with a selective transition from one portion of the speech user interface to another or invoking a voice messaging function being based upon the outcome of the evaluation performed in step b), wherein the steps of controlling transition and/or invoking are facilitated through an interpreter (134), the a speech user interface structure (138), a voice messaging function library (132), and/or a recognition command generator (142). 25 30

5. The method according to any of claims 1 to 4, wherein

selectively-interruptable prompts and messages to the subscriber are issued; wherein the step controlling the issuance is facilitated through the interpreter (134), the user interface structure (136, 138, 140), at least one voice messaging function within the voice messaging function library (132), and a phrase expander (232); wherein the phrase expander (232) is preferably responsive to signals issued by a DTMF processor (224) and the speech and logging supervisor (230) and is coupled to the speech recognizer (222). 40 45 50

6. The method according to any of claims 1 to 5, wherein

one or more portions of a DTMF user interface are made available in parallel with the speech user interface, and the DTMF user interface 55

works as a back-up in situations of repeated speech recognition failures; wherein the interpreter (134), the voice messaging function library (132) and a DTMF user interface structure (136) facilitate these features; and/or speech-responsive voice messaging system (SRVMS) performance information is selectively generated and analysed; wherein the generation and the analysation are facilitated by a logging unit (146), a reporting unit (148), and the speech and logging supervisor (230).

7. A speech-responsive voice messaging system (10) comprising a system control unit (100), a disk and voice Input/Output (I/O) control unit (160), a data storage unit (170) upon which a database directory entry and a mailbox for each subscriber reside, at least one digital line card (DLC) (180), a telephony interface controller (TIC) (185) corresponding to each DLC (180), and/or a system manager's terminal (SMT) (250); wherein the elements of the speech responsive voice messaging system (10) are selectively coupled via first control bus (260) and a first data base (262) and wherein each TIC (185) is provided to be coupled to a Central Office (CO) switch (20); wherein the system control unit (100) serves to manage the overall operation of the SRVMS (10), in accordance with system parameter settings received via the SMT (250) and the DLC (180) servers to exchange voice data with the CO switch (20), to process DTMF signals, and to perform speech recognition and logging operations under the direction of the system control unit (100). 35

8. The system of claim 7, wherein

the system control unit (100) comprises at least one or more of the following components a bus and Direct Memory Access (DMA) controller (110), a processing unit (120), and a memory (130) in which a voice messaging (VM) function library (132), an interpreter (134), a DTMF user interface structure (136), a speech user interface structure (138), and ambiguity resolution user interface structure (140), a recognition command generator (142), a recognition result processor (144), a logging unit (146), a reporting unit (148), a Speech Parameter Block (SPAB) library (150), and a call statistic library (152); wherein the bus and DMA controller (110), the processing unit (120), and each element within the memory (130) is coupled via an internal bus (270), the bus and DMA controller (110) is further coupled to the first data and control busses (260, 262), the SMT (250), and to the reporting system (12); wherein the coupling between the bus and DMA controller (110) and the reporting system (12) preferably

includes multiple lines; and/or
 the DLC (180) comprises a DLC bus controller (190), a DLC processing unit (200), a Coder/Decoder (CODEC) (210), and a DLC memory (220); wherein the DLC memory (220) 5
 comprises a speech recognizer (222), a DTMF processor (224), a template library (226), a logging buffer (228), a speech and logging supervisor (230), a phrase expander (232), an auto-response library (234), a Pulse Code Modulation (PCM) data buffer (236), and a signal conditioner (238); wherein each element within the
 DLC memory (220) is coupled to the DLC bus controller (190) and the DLC processing unit 10
 (200) via a second data bus (280), and wherein the DLC bus controller (190) is coupled to the
 DLC processing unit (200) via a second control bus (282) and to the first data and control
 busses (260, 262); and wherein the CODEC 15
 (210) is coupled to the PCM data buffer (236)
 and the DLC bus controller (190) for effecting
 DMA-type operations between the PCM data
 buffer (236) and the TIC (185).

9. Method for speech-responsive voice messaging 25
 comprising the steps of:

a) receiving an incoming call notification from a
 DLC bus controller (190) in an interpreter 30
 (134);
 b) transferring a reference to an initial Speech
 Parameter Block (SPAB) (300) to a recognition
 command generator (142) from the interpreter
 (134) in response to the call notification;
 c) selection (500) of the initial SPAB (300) by 35
 the recognition command generator (142);
 d) the recognition command generator (142)
 retrieves (502) at least one recognizer parameter
 within the selected SPAB (300);
 e) the recognition command generator (142) 40
 issues (504) a recognition parameter directive
 to the speech and logging supervisor (230);
 f) the recognition command generator (142)
 issues (506) a recognition request;
 g) a voice messaging function within the VM 45
 function library (132) selects a current prompt
 and issues (508) a prompt notification to a
 phrase expander (232);
 h) the phrase expander (232) issues the current
 prompt to the subscriber via the PCM data 50
 buffer (236), the CODEC (210) and the TIC
 (185);
 i) the recognition result processor (144)
 retrieves (510) the candidate result;
 k) the recognition result processor (144) evaluates 55
 (512) the quality of the returned candidate
 results;
 preferably the recognition result processor

(144) returns a value to the interpreter (134)
 that indicates the outcome of the evaluation;
 l) the interpreter (134) determines (514)
 whether recognition is to be repeated;
 m) repeating of the steps f) to l) and counting
 the repetitions, preferably with variations of
 step g), as long as a predetermined repeat
 count has not been exceeded;
 n) if the repeat count has been exceeded, the
 interpreter (134) transitions to the DTMF user
 interface and the method is ended;
 o) if the repeat count has not been exceeded,
 the interpreter (134) determines (520) whether
 recognition confirmation is required wherein
 recognition confirmation is required; if confir-
 mation is not required, the method is continued
 with step r); if confirmation is required, the
 method is continued with step p);
 p) the interpreter (134) selects a position or
 location within the ambiguity resolution user
 interface structure (140) and transfers a refer-
 ence to a confirmation SPAB (300) to the rec-
 ognition command generator (142);
 q) the interpreter determines whether the con-
 firmation was successful; if the confirmation
 was not successful, the method returns to step
 f); if the confirmation was successful the
 method continues with step r);
 r) the interpreter (134) transfers (530) control to
 a mapping function that maps the best candi-
 date word ID to a digit sequence;
 s) the interpreter (134) determines (532)
 whether the mapped digit sequence corre-
 sponds to a speech user interface navigation
 operation; if the mapped digit sequence corre-
 sponds to a speech UI navigation operation,
 the method is continued with step t), if not, with
 step o);
 t) the interpreter (134) selects (534) a position
 or location within the speech user interface; the
 method is continued with step w);
 o) the interpreter transfers (536) control to a
 voice messaging function that corresponds to
 the mapped digit sequence;
 v) the interpreter (134) updates (538) a com-
 munication count;
 w) the interpreter (134) determines (540)
 whether the subscriber's call is complete; if the
 subscriber's call is complete, the method ends,
 otherwise, the method is continued with step
 c).

10. Method for confirming a speech recognition result
 comprising the steps of:

a) retrieving (700) confirmation SPAB data;
 b) setting (702) of recognizer parameters;
 c) determinating a current confirmation prompt

and issuing a confirmation prompt notification;

d) issuing a recognition request;

e) retrieving (706) of recognition results;

f) evaluating the retrieved candidate results;

g) determinating (712) if the confirmation result was good; if the confirmation result was good, the method is continued with step h), otherwise with step k);

h) setting (718) of a successful confirmation status indicator; the method is continued with step m);

5 i) determinating (714) whether reconfirmation is required; if reconfirmation is required, the method preferably continues with step d); if reconfirmation is not required, the method continues with step l);

10 j) setting (716) an unsuccessful confirmation status;

15 k) restoring (720) the recognizer parameters and ending of the method.

20

11. Method for evaluating a speech recognition result comprising the steps of:

a) determinating (600) whether the candidate result set indicates that an unrecoverable error or a time out condition has occurred; if so, the method is continued with step l), if not, the method is continued with step b);

b) selecting (620) a (next) candidate result;

c) comparing (622) candidate results with current quality thresholds;

30 d) determinating (624) the candidate quality;

e) determinating (626) whether another candidate result requires consideration; if so, the method continues with step b), if not, the method continues with step f);

f) determinating (630) whether at least one candidate result has been designated as good; if so, the method continues with step k), if not, the method continues with step g);

35 g) determinating (640) whether a questionable candidate result is present; if so, the method continues with step h), if not, the method continues with step l);

40 h) setting (642) a confirmation required status indicator;

i) returning one of the good candidates; ending of the method;

k) setting (638) a good result status; the method is continued with step i);

45 l) setting (602) a bad result status;

m) incrementing (604) a repeat count;

n) determinating (606) whether a predetermined limit is exceeded; if so, the method is continued with step p); if not, the method is continued with step o);

50 o) setting (608) of a repeat status; and ending

5 of the method;

p) setting (610) a limit exceeded status; and ending of the method;

wherein preferably in step f), "k" is replaced by "q" and in step k), "l" is replaced by "t" with the additional steps of:

q) determinating (632) whether there are multiple good candidates; if so, the method is continued with step r), if not, the method is continued with step k);

r) examining (634) of the score differences between each good candidate result;

s) determinating (636) whether a difference threshold is exceeded; if so, the method is continued with step k), if the difference threshold is not exceeded, the method is continued with step h);

15 t) returning (644) of the best candidate; and ending of the method.

20

12. The method according to claim 10, wherein step f) is performed according to the method of claim 11.

25

13. A method for creating Customer Data Records comprising the steps of:

a) establishing (1000) of reporting conditions, preferably by retrieving the reporting parameters;

b) determinating (1002) whether a recognition-level CDR is to be generated; if so, the method continues with step c), if no recognition level CDR is to be generated, the method continues with step e);

c) recognition results are monitored (1004);

d) generation (1006) and/or update of recognition statistics;

e) determinating whether the current communication is complete; if so, the method is continued with step f); if the communication is not completed, the method is continued with step b);

f) determinating (1010) whether a communication level CDR is to be generated; if so, the method continues with step g), if no, communication level CDR is to be generated the method continues with step h);

30 g) generating (1012) of communication statistics;

h) determinating (1014) whether the call is complete; if not, the method continues with step b), if the call is complete, the method continues with step i);

i) determinating (1016) whether a call level CDR is to be generated; if so, the method continues with step k), if no call level is to be generated the method continues with step l);

35 k) generating (1018) of call statistics;

i) generating (1020) of each required CDR; and
ending of the method;
wherein:

the customer data records are preferably 5
generated in a reporting unit (148);
wherein the reporting unit (148) operates
transparently during a call, monitoring the
operation of the interpreter (134), the rec- 10
ognition command generator (142), and
the recognition result processor (144) to
track system performance and generate
CDRs in accordance with the reporting
parameters specified in each SPAB (300); 15
and wherein the reporting unit (148) main-
tains a set of statistics within the call statis-
tic library (152) for the generation of each
CDR.

14. A method of utterance logging comprising the steps 20
of:

- a) examining (800) logging parameters to
determine whether utterance logging is
required during a current call; if not the method
ends; if utterance logging is required, the
method is continued with step b); 25
- b) establishing (802) of logging conditions;
- c) determinating (804) whether the next recog-
nition event is to be logged; if not, the method is
continued with step k), if the next recognition
event is to be logged, the method continues
with step d);
- d) setting (806) a set of at least one sampling
parameters; 30
- e) determinating (808) whether the quality eval-
uation for the current candidate results has
been completed; if not, the method preferably
remains at step e), if the quality evaluation has
been completed the method continues with
step f);
- f) determinating (810) whether any criteria
specified in the logging parameters are
matched; if not, the method is continued with
step k); if any criteria specified in the logging
parameters are matched, the method is contin- 35
- ued with step g);
- g) saving (812) of the utterance;
- h) examining of the current logging parameters
and determination whether the saved utterance
should be retained for later use; if the utterance
should not be retained, the method continues
with step k), otherwise with step i);
- i) updating (816) of the logging buffer storage
location referenced to a next available location; 50
- k) determinating (818) whether the current call
is complete; if not, the method preferably con-
tinues with step c); if the call is complete the 55

method continues with step l);
l) generating (820) call header information;
m) transferring (822) the logging data to a data
storage device; end of the method;
wherein:

preferably the method is performed by
means of a logging unit (146) and a
speech and logging supervisor (230)
which control the selective logging of utter-
ances in accordance with the logging
parameters specified in each of a group of
SPABs (300);
preferably the logging unit (146) operates
transparently during a call, monitoring the
operation of the interpreter (134), the rec-
ognition command generator (142) and the
recognition result processor (144);
the call header information comprises pref-
erably a mailbox ID, and possibly a refer-
ence to a Customer Data Record (CDR),
wherein the customer data records are
preferably generated by the method of
claim 13.

15. The method according to any of claims 1 to 6 or 9,
comprising the method of claim 11 and/or the
method of claim 10 or 12 and/or the method of
claim 14 and/or the method of claim 13.




Fig. 1

USER INTERFACE




Fig. 2




Fig. 3




Fig. 4A




Fig. 4B




Fig. 5




Fig. 6




Fig. 7




Fig. 8





Fig. 9A






Fig. 10

THIS PAGE BLANK (USPTO)